REMARKS

This Response is intended to fully respond to the Office Action dated June 23, 2003. In that Office Action, claims 1-20 were examined and rejected. More specifically, claims 1-4, 11-14 and 20 were rejected under 35 U.S.C. § 103(a) as being unpatentable over US Patent No. 6,289,360 issued to Kolodner et al., (hereinafter "Kolodner") in view of 5,787,447 issued to Smithline et al. (hereinafter "Smithline") while claims 5-10 and 15-19 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Kolodner, by itself.

Reexamination and reconsideration are respectfully requested in light of these remarks. Claims 1-20 are pending in the present application.

Claim Rejections – 35 U.S.C. § 103(a)

Claims 1-4, 11-14 and 20 were rejected under 35 U.S.C. § 103(a) as being unpatentable over US Patent No. 6,289,360 issued to Kolodner et al., (hereinafter "Kolodner") in view of 5,787,447 issued to Smithline et al. (hereinafter "Smithline"). Moreover, claims 5-10 and 15-19 were rejected under 35 U.S.C. § 103(a) as being unpatentable over US Patent No. 6,289,360 issued to Kolodner, by itself.

Applicant respectfully traverses the section 103 rejections. The Examiner has failed to substantiate a prima facie case of obviousness because one of the requirements of a prima facie case is absent. Indeed, such a prima facie case can only be met when **all** of the following requirements are met: (1) there must be some suggestion or motivation in the references themselves (or in the knowledge available to those skilled in the art) to combine the references; (2) there must be a reasonable expectation of success; and (3) the combined references must teach or suggest all the claim limitations. See MPEP §§ 706.02(j) and 2143. In this case, none of the requirements are met. For instance, the combination of Kolodner and Smithline does not teach all of the claim limitations, e.g., the dedication of each heap of memory to one of a plurality of processing units the explicit or implicit synchronization of the plurality of processing units during the garbage collection process.

The present invention relates to a method of using multiple processors in its garbage collection process. In using these multiple processors, the memory is divided into separate

heaps, and each heap is dedicated to a processor. The separate processors perform several phases of garbage collection on their own dedicated heaps. Importantly, the processors are synchronized because each processor may access objects located in other heaps. Thus, each phase must be completed by all processors before starting another phase. Furthermore, in order to accommodate multiple processors operating on the multiple heaps of memory, the application process is paused during the garbage collection phase.

In the present Office Action, the Examiner has cited Kolodner as disclosing garbage collection that is performed in phases, such as mark and sweep. Applicants agree that Kolodner describes mark and sweep phases in its particular garbage collection process.

Further however, the Examiner has cited Kolodner as disclosing the use of multiple processors, wherein "each processor performs each of the phases on the heap dedicated to the processor using a garbage collection thread executing on the processor (collector threads to force synchronization process: col. 3, lines 25-45 and col. 5, lines 4-14)." Applicant respectfully disagrees with this assertion. Kolodner mentions multiprocessor systems and in fact, describes problems with multiprocessor systems as they apply to concurrent garbage collection systems. However, Kolodner never states that different processors may be dedicated to separate heaps. Indeed, Kolodner does not necessarily divide the memory into heaps (as acknowledged by the Examiner) such that there cannot be an inherent dedication of heaps to processors.

Moreover, the "synchronization process" discussed in Kolodner is the coordination of a mutator thread and a collector thread in a **concurrent** garbage collection process. The mutator thread in Kolodner "is the running program ... [that] changes object graphs." Thus, the mutator thread must be synchronized in concurrent garbage collection processes with the collector thread in order to prevent newly allocated memory to be inadvertently treated as unreachable memory. Indeed, Kolodner describes a new method that reduces the dependency on this required coordination between the mutator and the collector threads. Consequently, the only interpretation of the Kolodner, with respect to multiple processors, relates to the synchronization of the mutator thread on one processor and the collector thread on another, unless impermissible hindsight is used.

With respect to the present application, the process described is not a concurrent garbage

10

collection process. The "mutator" of the present invention is paused during the entire garbage collection process (See Fig. 5). Therefore, the discussion of synchronization in Kolodner does not, in any way, equate to the synchronization of a plurality of processors as described in the present invention. Moreover, since Kolodner does not divide the memory in to heaps (as acknowledged by the Examiner), there is no reason to (a) use separate collector threads or (b) synchronize separate collector threads.

The Examiner further states that Kolodner discloses a process involving "synchronizing the processors so that all processors have completed the preceding phase prior to each processor beginning the next phase (the beginning phase and the ending phase of the mark-sweep cycle and the synchronization process between the mark-sweep phases: abstract, col. 2, lines 57-67; also col. 3, lines 25-45 and col. 5, lines 44-61)." Applicant also respectfully disagrees with this assertion. As stated above, Kolodner does not describe the synchronization of multiple processors as claimed in the present application. More particularly, Kolodner does not describe the synchronization of multiple collector threads thereby allowing access of any one processor into other, non-dedicated heaps. Since Kolodner does not dedicate heaps of memory to separate processors (or disclose such a possible configuration), it simply does not disclose the synchronization process claimed in the present application. Moreover, as stated above, the synchronization described in Kolodner is between the mutator and the collector threads. There is absolutely no discussion of any reason to have multiple collector threads, no discussion of how such an imaginary plurality of collector threads would operate, and clearly no discussion of how or why the collector threads would be synchronized.

Smithline does not satisfy the inadequacies of Kolodner. Smithline describes the use of memory heaps and the balancing of the heaps for management benefits. However, Smithline does not dedicate heaps to processors. Indeed, Smithline does not contemplate the use of more than one processor. (See col. 3, lines 15-19). Consequently, the combination of Smithline and Kolodner does not disclose the dedication of processors to heaps in order to provide parallel garbage collection.

Given that the combined references do not teach or suggest all the claim limitations of the present invention, all claims are believed to be allowable over the prior art. That is, because

11

Kolodner does not disclose multiple heaps, each dedicated to a separate collector thread, and since Smithline does not disclose the dedication of collector threads and/or processors to memory heaps, claims 1, 5, 10, 11, 15 and 20 are not obvious in light of Kolodner and Smithline. Therefore these claims are believed to be allowable. Similarly, all claims depending from those claims are also allowable.

Conclusion

As originally filed, the present application included 20 claims, 6 of which were independent. No amendments have been made and no new claims have been added. A Petition for a one-month extension of time is included with this Response, together with the fee of $110. It is believed that no further fees are due with this Response. However, the Commissioner is hereby authorized to charge any deficiencies or credit any overpayment with respect to this patent application to deposit account number 13-2725.
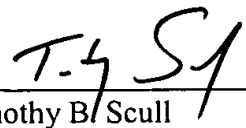
In light of the above remarks, it is believed that the application is now in condition for allowance, and such action is respectfully requested. Should any additional issues need to be resolved, the Examiner is requested to telephone the undersigned to attempt to resolve those issues.

Respectfully submitted,

Merchant & Gould, P.C.
3100 Norwest Center
Minneapolis, Minnesota 55402-4131
303.357.1648

Date: October 21, 2003

27488
PATENT TRADEMARK OFFICE

Timothy B. Scull
Reg. No. 42,137
MERCHANT & GOULD P.C.
P.O. Box 2903
Minneapolis, MN 55402-0903
(303) 357-1648

12